

DESIGN OF A HIGH PERFORMANCE CONCURRENT  
SIMULATION SYSTEM

FINAL TECHNICAL REPORT

VERNON REGO AND VAIDY SUNDERAM

SEPTEMBER 1, 1997

U.S. ARMY RESEARCH OFFICE

CONTRACT/GRANT NUMBER  
DAAH04-93-G-0045

PURDUE UNIVERSITY  
WEST LAFAYETTE, IN 47907

APPROVED FOR PUBLIC RELEASE:

DISTRIBUTION UNLIMITED

**DTIC QUALITY INSPECTED 2**

THE VIEWS, OPINIONS, AND/OR FINDINGS CONTAINED IN THIS  
REPORT ARE THOSE OF THE AUTHOR(S) AND SHOULD NOT BE  
CONSIDERED AS AN OFFICIAL DEPARTMENT OF THE ARMY  
POSITION, POLICY, OR DECISION, UNLESS SO DESIGNATED BY  
OTHER DOCUMENTATION.

19971203 090

<b>REPORT DOCUMENTATION PAGE</b>			Form Approved OMB NO. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimates or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 1, 1997		3. REPORT TYPE AND DATES COVERED Final Report 3/1/93 - 8/31/97
4. TITLE AND SUBTITLE Design of a High Performance Concurrent Simulation System			5. FUNDING NUMBERS DAAH04-93-G-0045	
6. AUTHOR(S) Vernon Rego and Vaidy Sunderam				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Computer Science Purdue University West Lafayette, IN 47907			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSORING / MONITORING AGENCY REPORT NUMBER ARO 3160.23-MA-SDI	
11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT  Approved for public release; distribution unlimited.			12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) We investigate methodologies and tools for the deployment of massively parallel simulation-oriented computations on a variety of hardware platforms, particularly heterogeneous networks of workstations. The underlying ideas are based in part on the prize-winning performance of the Eclipse parallel simulation toolkit in the 1992 Gordon Bell Prize competition. The primary motivation is that the replication of computations involving statistical sampling be deployable through a transparent use of sequential codes. A major goal of this research is the design and experimental development of a high-performance simulation environment which enables large-scale computations in different problem domains. Besides execution performance, the research emphasizes low-effort and rapid simulator development, a feature invaluable in experimental research.				
14. SUBJECT TERMS simulation, parallel, heterogeneous networks, threads, objects, process, stochastic			15. NUMBER OF PAGES	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

REPORT DOCUMENTATION PAGE (SF298)  
(Continuation Sheet)

**Final Report (September 1, 1997)**

In this project we focused on the design of a high performance concurrent simulation system. For this, we took an experimental software systems approach, based on our PVM/Conch experience, to study different aspects of support for high-performance distributed simulation on heterogeneous networks. We designed new algorithms and methods for deploying distributed simulations on heterogeneous processor clusters. We built a layered software prototypes, incorporating functionality for heterogeneous networked computing with threads and objects. We placed major emphasis on failure-resilient execution in one of these subsystems (i.e., EclIPSe).

Our research required us to delve into a number of related areas for which synthesis is not well understood, particularly with threads-oriented distributed computation. These areas include heterogeneous network computing, efficient networking (lightweight multicast protocols), distributed visualization, distributed naming and location, parallel stochastic and deterministic simulation, multithreaded operating systems, statistical computing and fault-tolerance.

Our research methodology and software enables users to create complex simulation models rapidly, because of the "process" orientation. Dynamic simulation units are viewed as processes – implemented as user-space threads, and static simulation units may be viewed as objects. Distributed simulations proceed both optimistically and adaptively. Our solution methodology is unique in two senses. First, there is virtually no commercial software available for (optimistic) distributed simulation. Second, our ParaSol system provides a simple API due to its process-orientation (i.e., its well-accepted that process-oriented models tend to be the easiest to design) that is supported by migratable threads. The thread migration support makes the system unique. To the best of our knowledge, this is the first known process-oriented parallel simulation tool.

We focused mainly on the ParaSol simulation object library for distributed simulation and implemented (a). the ParaSol kernel, based on an optimistic simulation protocol, and (b). a domain layer for manufacturing applications (queueing), which runs on top of the kernel. The kernel is C++ based, and is supported by our Ariadne portable threads library. We are currently investigating the layering of other application domains on top of the ParaSol kernel, potentially domains of interest to researchers at the ARL. Some projects under way include Professor Hisao Nakanishi's particle domain for polymer physics (Purdue), Dr. Danny Rintoul's materials science simulation domain (Princeton Univ. Material Science Lab.), Prof. Ferydoon Family's colloidal coagulation and stability models using the NVT method (Emory Univ.), and Professor Bruce Schmeiser's queueing domain for manufacturing systems (Purdue).

REPORT DOCUMENTATION PAGE (SF298)  
(Continuation Sheet)

**Publications**

1. K. Chung, J. Sang and V. Rego, A Performance Comparison of Event Calendar Algorithms, *Software - Practice & Experience*, Vol. 23(10), pp. 1107-1138, October 1993.
2. K. Chung, J. Sang and V. Rego, Sol.es: An object-oriented platform for event-scheduled simulations, *Proc. of 1993 Summer simulation Conference*, pp. 972-977, Jul 19-21, Boston, MA 1993.
3. J. Sang, E. Mascarenhas and V. Rego, Process Mobility in Distributed Memory Simulation Systems, *Proc. of Winter Simulation Conference*, Los Angeles, December 1993.
4. K. Chung, A Concurrent Composite Computational Model for Stochastic Simulation, *Ph.D. Dissertation, Dept. of Computer Sciences, Purdue University, July 1993*.
5. B. Schmidt, V. S. Sunderam, Empirical Analysis of Overheads in Cluster Environments, *Concurrency - Practice and Experience*, 1993.
6. V. S. Sunderam, DCL: Protocols and Primitives for Distributed and Concurrent Computing in Networked Environments, *International Conference on Computers and Communication (ICCI 93)*, May 1993.
7. B. Topol, Second Generation Heterogeneous Concurrent Computing with Conch, *M. S. Thesis, Dept. of Math & CS, Emory University, May 1993*.
8. K. Chung, J. Sang and V. Rego, Dynamic Performance Profiles of Simulation Calendars, *26th Annual Simulation Symposium, Simulation MultiConference*, Washington, D.C., Mar 31, 1993.
9. J. Sang, K. Chung and V. Rego), Computational Schemes for Efficient Simulation of Service Disciplines, *26th Annual Simulation Symposium, Simulation MultiConference*, Washington, D.C., Mar 31, 1993.
10. J. Sang, F. Knop, V. Rego, J. Lee and C. King, The Xthreads library: Design, Implementation, and Applications, *Proc. of The 17th Annual International Computer Software and Applications Conference (COMPSAC '93)*, Phoenix, Arizona, November 1993.
11. J. Sang, K. Chung and V. Rego, Design and implementation of a Simulation Library using Lightweight Processes, *Proc. of Usenix Summer Technical Conference*, Cincinnati, OH, 1993.
12. Experiments with Program Unification on the Cray Y/MP (L. Chuang, V. Rego, and A. Mathur), *Concurrency, Practice & Experience*, Vol. 6, No. 1, pp. 33-54, February 1994.
13. A Simulation Testbed based on Lightweight Processes, (J. Sang, K. Chung and V. Rego), *Software, Practice & Experience*, Vol. 24, No. 5, pp. 485-505, May 1994.
14. Efficient Algorithms for Simulating Service Disciplines, (J. Sang, K. Chung and V. Rego), *Simulation, Practice & Theory*, Vol. 1, pp. 223-244, 1994.
15. Recent Initiatives in Heterogeneous Parallel Computing, (V. Sunderam), *Parallel Computing: Technology and Practice*, Eds. J. P. Gray and F. Naghdy, pp. 1-16, IOS Press, Amsterdam, November 1994.
16. Heterogeneous Concurrent Computing with Exportable Services, (V. Sunderam), *Environments and Tools for Parallel Scientific Computing*, Eds. J. J. Dongarra and B. Tourancheau, pp. 142-151, SIAM Press, May 1994.
17. General Purpose Parallel Computing With PVM, (V. Sunderam), *Parallel Processing Comes of Age*, Ed. Francis Wray, Chapman Hall, London, June 1994.
18. *EcliPSe*: A System for Fault-Tolerant Replicative Computations (F. Knop, V. Rego and V. Sunderam), *IEEE/USP International Workshop on High Performance Computing (WHPC '94)*, Sao Paulo, Brazil, March 28-30, 1994.

REPORT DOCUMENTATION PAGE (SF298)  
(Continuation Sheet)

19. Failure-resilient Computations in the *EclIPSe* System (F. Knop, A. Ferrari, V. Rego and V. Sunderam), *Proceedings of the 23rd Annual and International Conference on Parallel Processing*, August 15-19, 1994.
20. An Introduction to Fault-tolerant parallel simulation with *EclIPSe*, (F. Knop, E. Mascarenhas, V. Rego and V. Sunderam), *Proc. Winter Simulation Conference*, Florida, Dec. 13-16, 1994.
21. Thread Migration on Heterogeneous Systems via Compile-Time Transformations, (J. Sang, G. Peters and V. Rego), *Proc. International Conference on Parallel and Distributed Systems*, National Chiao Tung University, Hsinchu, Taiwan, ROC, Dec. 19-21, 1994.
22. E. Mascarenhas, V. rego and J. Sang, DISplay: A System for Visual-Interaction in Distributed Simulations, *Proceedings of The Winter Simulation Conference*, Washington, D.C., December 3-6, 1995.
23. F. Knop and V. Rego, Parallel Cluster Labeling on a Network of Workstations, *The 13th Brazilian Symposium on Computer Networks, SBRC '95*, Belo Horizonte, Brazil, May 22-26, 1995.
24. E. Mascarenhas, F. Knop and V. Rego, ParaSol: A multithreaded System for Parallel Simulation based on Mobile Threads, *Proceedings of The Winter Simulation Conference*, Washington, D.C., December 3-6, 1995.
25. H. Nakanishi, V. Rego and V. Sunderam, On the Effectiveness of Superconcurrent Computations on Heterogeneous Networks, *Journal of Parallel & Distributed Computing*, Vol. 24, No. 2, pp. 177-190,
26. F. Knop, E. Mascarenhas, V. Rego and and V. Sunderam, Concurrent and fail-safe replicated simulations on heterogeneous networks: An Introduction to *EclIPSe*, *Simulation, Practice & Theory*, Vol. 3, pp. 121-146, November 15, 1995.
27. E. Mascarenhas and V. Rego, An Architecture for Visualization and User Interaction in Parallel Environments *Computers & Graphics*, Volume 19, No. 5, pp. 739-753, September/October 1995.
28. F. Knop, E. Mascarenhas and V. Rego, A Parallel GPSS based on the ParaSol Paradigm: Design and Implementation, *Proceedings of The Winter Simulation Conference*, San Diego, CA, December 8-11, 1996.
29. J. Sang, E. Mascarenhas and V. Rego, Mobile-Process Based Parallel Simulation, *Journal of Parallel & Distributed Computing*, Vol. 33, No. 1, pp. 12-23, February 1996.
30. E. Mascarenhas and V. Rego, Ariadne: Architecture of a Portable Threads System supporting Thread Migration, *Software, Practice & Experience*, Vol. 26, No. 3, pp. 327-356, March 1996.
31. F. Knop, V. Rego and V. Sunderam, Fail-Safe Concurrency in the *EclIPSe* System, *Concurrency, Practice & Experience*, Vol. 8, No. 4, pp. 283-312, May 1996.
32. I. Zoraja, H. Hellwagner and V. Sunderam, SCI-PVM: Parallel and Distributed Computing on SCI Workstation Clusters, Technical Report, Dept. of Math and Computer Science, Emory University, 1996.
33. A. Krantz and V. Sunderam, Client Server Computing on Message Passing Systems: Experiences with the RPC Model in PVM, Technical Report, Dept. of Math and Computer Science, Emory University, 1996.
34. T. Goddard and V. Sunderam, WebVector: Agents with URLs, Technical Report, Dept. of Math and Computer Science, Emory University, 1996.
35. F. Knop, V. Rego and V. Sunderam, Fail-Safe Concurrency in the *EclIPSe* System, *Concurrency, Practice & Experience*, Vol. 8, No. 4, pp. 283-312, May 1996.

REPORT DOCUMENTATION PAGE (SF298)  
(Continuation Sheet)

36. J.-C. Gomez, E. Mascarenhas, and V. Rego, The CLAM Approach to Multithreaded Communication on Shared-memory Multiprocessors: Design and Experiments, Technical Report, Dept. of Computer Sciences, Purdue University, 1996.
37. J.-C. Gomez, V. Rego, and V. Sunderam, Efficient MultiThreaded User-Space Transport for Network Computing: Design and Test of the TRAP Protocol *Journal of Parallel & Distributed Computing*, (accepted for publication), Technical Report, Computer Sciences, Purdue University, 1996.
38. F. Knop, Software Architectures for Fault-Tolerant Replications and Multithreaded Decompositions: Experiments with Practical Parallel Simulation, Ph.D. Dissertation, Dept. of Computer Sciences, Purdue University, August 1996.
39. E. Mascarenhas, A System for Multithreaded Parallel Simulation and Computation with Migrant Threads and Objects, Ph.D. Dissertation, Dept. of Computer Sciences, Purdue University, August 1996.

### Personnel

Scientific personnel supported include V. Rego, V. Sunderam, A. Krantz, K.-H. Chung, J. Sang, F. Knop, E. Mascarenhas, J.-C. Gomez.

### Scientific Progress

This project has demonstrated several things: process-oriented parallel simulation is feasible, user-space protocols are very efficient, and a variety of services can be provided based on URLs. Our experimental research has brought significant gains in knowledge, particularly with respect to threads and thread migration, the use of threads in the design of kernels for optimistic simulators, experiences with the use of statistics-based protocol adaptation. Besides their use in protocols for simulation, we have found threads to be highly useful in the design of user level protocols in networks. Our experimental work on user level implementation of a variety of protocols is now continuing. We expect the communication protocols to be very useful in improving the performance of distributed applications.

REPORT DOCUMENTATION PAGE (SF298)  
(Continuation Sheet)

**Awards/Honors**

For each of the past several years, researchers on this project (students of Rego) have been given annual awards for the above work.

1. 1992: Rego and Sunderam are awarded the 1992 IEEE/Gordon Bell Prize for practical and large-scale parallel computing.
2. 1994: Janche Sang (Ph.D. Student) is awarded the 1994 Maurice Halstead Award in Software Engineering.
3. 1993: The first student from the group (with research fully supported by the Taiwan DoD during his stay at Purdue, and who graduated in 1993) was promoted to the rank of General in the Taiwanese Army.
4. 1995: Felipe Knop (Ph.D. Student) is awarded the 1995 Maurice Halstead Award in Software Engineering.
5. 1996: Edward Mascarenhas (Ph.D. Student) is awarded the 1996 Maurice Halstead Award in Software Engineering.
6. 1997: J.-C. Gomez (Ph.D. Student) is awarded the 1997 Maurice Halstead Award in Software Engineering.